Choosing DACs for Direct Digital Synthesis

by David Buchanan

INTRODUCTION

Direct Digital Synthesis (DDS) is a technique for deriving, under digital control, an analog frequency source from a single reference clock frequency. This technique provides high frequency accuracy; temperature and time stability; wideband tuning; and very fast, phase continuous frequency tuning. The performance of synthesizers using this technique is often limited by the performance of available digital-to-analog converters (DACs). This application note explains the basic architecture of DDS and some of its advantages in system design, and outlines some of the performance characteristics designers should look for when choosing DACs for DDS applications. Performance tradeoffs are also explored, and recommendations are made for characterizing the DACs.

DDS BACKGROUND

A simplified block diagram of a direct digital synthesizer is shown in Figure 1. The synthesizer has two digital inputs: a frequency control word (Δ Phase) and a reference clock signal (f_c). The output of the synthesizer is an analog sine wave with frequency f_a. The relationship between f_c and f_a is determined as

$$f_a = \frac{\Delta\, PHASE}{2^N}\, f_c$$

Where N is the resolution of the frequency control word, Phase.

As illustrated in Figure 1, the circuit is easily divided into three blocks: a phase accumulator, a phase-to-sine converter, and a DAC. The first two blocks are both digital circuits. The phase accumulator is simply an adder with a programmable step size, Δ Phase, representing the phase step taken by the output waveform during each clock cycle. On each clock cycle, the phase accumulator output represents the phase of the output sine wave, with all zeros representing 0 radians, and all ones representing 2π. This signal is a digital ramp with a frequency equal to the output sine wave.
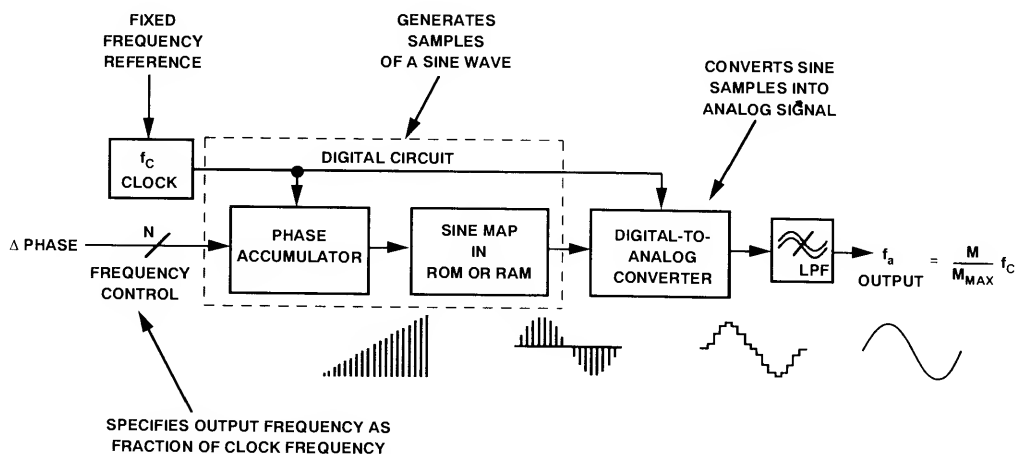

Figure 1. Block Diagram of DDS Generator

The phase-to-sine converter circuit takes the M most significant bits of the phase accumulator's output and provides an M-2 bit sine amplitude output; M-2 determines the resolution of the circuit to follow, usually a DAC. Truncation of the N-M least significant bits is necessary to reduce the complexity of the phase-to-sine conversion. This function may be performed by a look-up table stored in memory, or the sine value may be calculated from a digital algorithm to make a faster or smaller circuit.

The phase accumulator and phase-to-sine converter together form a DDS system with a digital output. The digital output is useful in many applications as a frequency reference (digital demodulation as an example), but most applications require a transformation of the digital sine wave into an analog frequency reference. This makes the digital to analog converter extremely important.

DDS PERFORMANCE CHARACTERISTICS

DDS has both advantages and disadvantages over other frequency synthesis techniques such as phased locked loops. While it is not the intent of this paper to explore fully the tradeoffs of choosing DDS architecture for a synthesizer, it will point out some of the more obvious ones and discuss the important performance characteristics.

A DDS system effectively provides a frequency reference that is a fraction of the clock input frequency. The DDS is digitally tuned by the Δ Phase input, usually controlled by a microcontroller or digital signal processor. Once the digital data are registered on board the phase accumulator, the controlling circuits are free to perform other functions in the system. The digital nature of the DDS eliminates the inconvenience associated with the "tweaking" of synthesizer designs that rely on analog component values to determine frequencies.

The frequency resolution is determined by N, the resolution of Δ Phase, as

$$\frac{1}{2^N} f_c$$

As an example, if the DDS clock reference, f_c, is 20 MHz, and N = 32, the frequency resolution of the synthesizer will be 4.66 millihertz (mHz). This is an advantage over phased locked loops, in which the reference frequency directly determines the frequency resolution and must be large enough to avoid large multiplication ratios. Most integrated solutions for DDS provide at least 24 bits of frequency resolution; some provide up to 48 bits. Many phase accumulators are designed so they can be cascaded to increase the frequency resolution.

DDS provides not only very accurate frequency resolution, but also provides a wide frequency range. As described above, the output frequency is determined as a fraction of the clock frequency reference. The lowest frequency is the smallest fraction of the clock that can be

used, or the resolution of the phase accumulator (see above). Nyquist's theory holds that the DAC can reproduce signals up to one half of this clock frequency. Thus the upper limit on the frequency range of the DDS system is determined by the maximum clock rate the synthesizer will support $f_{MAX} = (f_c/2)$. Technology advances over time have increased the available clock rates of both the digital and DAC portions of the DDS circuit.

DDS TECHNOLOGY TODAY

As in other digital circuits, the phase accumulator and phase-to-sine conversion circuit designs must be optimized for cost and power. The DAC design must also address power and cost concerns, but dynamic performance of the converter is of premium consideration.

There are CMOS digital devices available that provide DDS solutions for clock rates up to 100 MHz. A few bipolar devices cover clock frequencies up to 300 MHz, and there are also GaAs devices that provide digital solutions up to 1.4 GHz clock rates.

Most DACs used in DDS are bipolar, although a few GaAs DAC designs see service in the higher frequency applications. Designers prefer monolithic DACs to keep the cost of the converter reasonable. Some 12-bit monolithic DACs can clock up to 100 MSPS (AD9713B), while higher speed applications can take advantage of monolithic 10-bit designs that clock up to 400 MSPS (AD9720). Above 400 MSPS, there are a few 8-bit devices. As a later section of this paper will point out, resolution and speed of a DAC do not always determine the DAC's suitability for DDS applications.

ADVANTAGES AND DISADVANTAGES

While the available clock frequencies described above indicate that DDS circuits can generate output frequencies well into the UHF band, in practice the frequency range of a DDS system is limited by the real world characteristics of the DACs. This is because the resolution and performance (in terms of power, cost, ease of use) that can be practically realized in the DACs at higher clock frequencies are limited. Designers may find that a particular system design could benefit from a DDS which would generate a frequency reference directly, but the required clock rate dictates power, cost, and performance tradeoffs that are unacceptable.

Often the benefits of the DDS can be preserved by using a hybrid design approach. An example of this would be to use DDS at an intermediate frequency with a mixer and second frequency reference, or a PLL to translate the signal to the RF band.

Furthermore, as the output frequency increases, the number of amplitude samples for each sine cycle decreases, making it more difficult for the DAC to reproduce accurately the output sine wave. The accuracy of the analog sine wave is often described in terms of its spectral purity (see below). Each system application of DDS will have a limit on spectral purity. In many applications, the DDS system will meet the spectral purity specifications over only a portion of the available $f_c/2$ bandwidth, effectively limiting the frequency range of the synthesizer. As a rule of thumb, the output frequency should be limited to one fourth of the clock frequency to preserve a reasonable level of spectral purity.

Figure 2 illustrates another feature of DDS, its frequency tuning speed. If the phase input is changed, the phase accumulator instantaneously changes to the new frequency at the next clock cycle. Pipeline delays used in the digital circuits are usually the only limit to switching speed. Frequency transitions in DDS are also phase continuous. Figure 2 illustrates how this instantaneous, phase continuous frequency switching might compare to that of a phase locked loop. The PLL transition includes a frequency transition period ($t_{SETTLING}$, usually a few microseconds) and a frequency overshoot. Coherent analog frequency synthesizers (fixed oscillators, mixers, and filters) also have fast switching speed, but are not phase continuous and cannot be implemented easily.
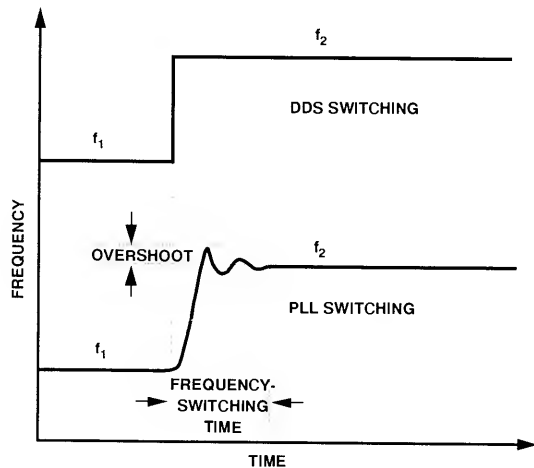

Figure 2. Comparison of DDS and PLL Frequency Switching Characteristics

DDS also has inherently low phase noise and drift. These characteristics are essentially those of the reference clock, f_c. In most DDS applications, a fixed crystal oscillator provides the reference frequency and therefore the phase noise and drift characteristics are excellent.

As mentioned earlier, the spectral purity of a DDS is often a limiting factor in its performance, and is determined by the dynamic performance of the DAC. Ideally, a synthesizer's output spectrum consists of a single frequency.

Since a DDS's output is based on a digital approximation of the sine wave, the theoretical output spectrum contains only the frequency of interest and a constant quantization noise of $q/(12)^{1/2}$ (where q is the weight of one LSB) spread evenly from dc to $f_c/2$. This ideal spectrum is illustrated in Figure 3 (which ignores phase noise). As with all sampled systems, the amplitude response of the output frequency is weighted as

$$A = \frac{sin\ (\pi\ f_a/f_c)}{(\pi\ f_a/f_c)}$$
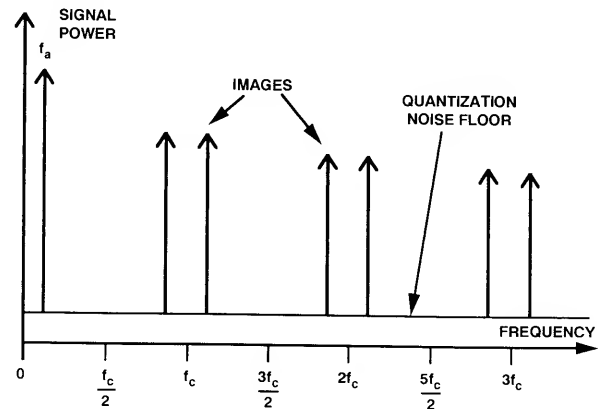
where A is the normalized output amplitude.

Figure 3. Ideal DDS Output Spectrum

This effect can be corrected by a digital, inverse (sine x)/x filter. Images of the Nyquist bandwidth (dc to $f_c/2$) will appear around each multiple of f_c, again weighted by the (sine x)/x function. These images are commonly removed with a low pass filter at the output of the DAC.

A more realistic illustration of a DDS output spectrum is shown in Figure 4. Here it becomes evident that transfer functions are not ideal. The additional signal content is generated by the digital-to-analog conversion process. The noise contribution is no longer uniform, and harmonics of the fundamental frequency and its images are created, along with other frequencies that have no obvious harmonic relationship.

In many applications, a synthesizer must meet a specified spurious free dynamic range (SFDR) over its output bandwidth. This specification defines the difference in power of the signal of interest and the worst case (highest) signal power of any other signal in the band of interest. This concept is illustrated in Figure 4 for the Nyquist bandwidth. Some other definitions of SFDR do not include direct harmonics, despite the fact the frequency band of interest may include them.
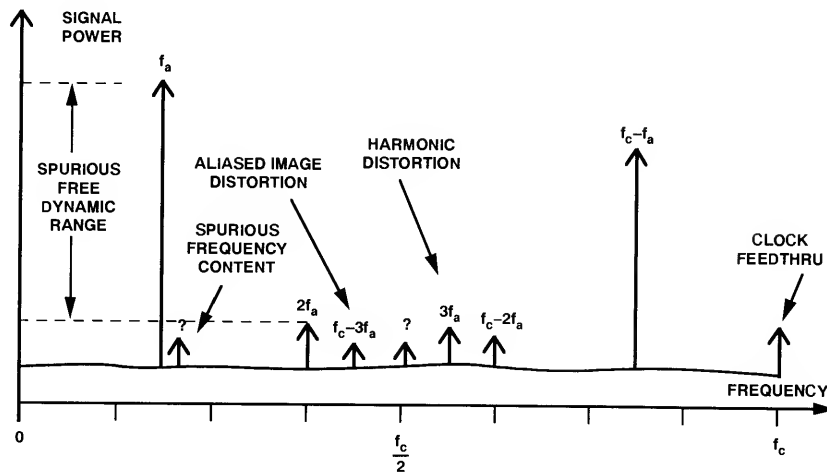
Figure 4. Real DDS Output Spectrum

It is difficult to generalize about the SFDR of direct digital synthesizers because of the number of variables involved. The resolution of the DAC, its ac characteristics, clock rate, and the circuit layout all come into play. DAC characteristics that contribute to these errors are discussed later. Even if the DAC contributions were constant, generalizations might still be misleading. For instance, the illustration in Figure 4 defines a SFDR for the Nyquist bandwidth; an application whose band of interest was limited from $f_c/8$ to $f_c/4$ would see an improvement in this specification. Applications also vary widely in their requirements for spectral purity. System designs which take advantage of the frequency accuracy and stability for channel selection often require a very high SFDR, and are thus restricted to smaller bandwidths. Systems that take advantage of the digital interface to implement IF modulators may be more forgiving in their spectral requirements, and therefore can use a wider bandwidth.

As the discussion of DACs below indicates, the only safe generalization about spectral purity is "each DDS application must be proven out and characterized in the lab." It is unfair, however, to prepare this application note without giving the reader some idea of the performance that can be achieved with available technology.

Ten- to twelve-bit DACs have been shown to provide up to 70 dB of SFDR in applications with clock rates less than 80 MHz and analog bandwidths less than a few MHz. Designs using higher clock rates have been limited to 8-bit DACs, and a SFDR of up to 45 dB is often quoted as a limit. Caution is advised in using these figures because they may not be applicable over wide analog bandwidths.

The modulation capabilities of a synthesizer are also important, and DDS is versatile in this respect. Frequency modulation is possible directly through the Δ phase data port. In fact, circuits ahead of the DDS could use the Δ phase input to specify both channel selection and FM modulation simultaneously.

Figure 5 shows a DDS block diagram modified to include digital control of the phase modulation (PM) and amplitude modulation (AM) of the waveform. Thus, a user could implement all three forms of modulation digitally with a highly integrated DDS.
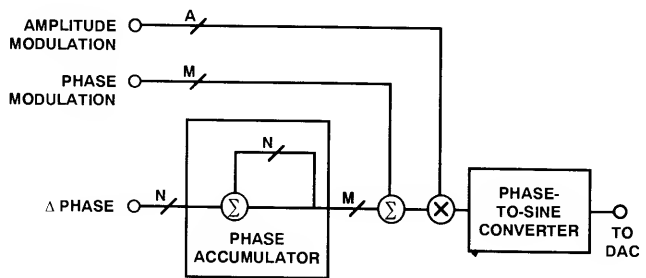

Figure 5. Modulated DDS

Amplitude modulation can also be implemented by modulating the reference to the DAC. Figure 6 shows a circuit using a second DAC to modulate the DDS DAC's reference under digital control. In such an application, the amplitude modulation data would have to be skewed in time to match any pipeline delays in the digital part of the DDS system. Many DAC designs integrate a reference and reference amplifier, but only a few provide enough bandwidth to support multiplying applications.
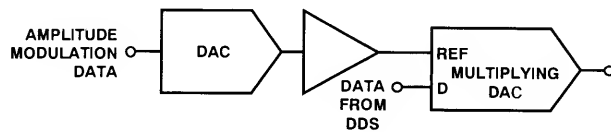
*Figure 6. Amplitude Modulation Using
a Multiplying DAC*

To this point, DDS has been described as a wideband, frequency agile synthesizer with digital tuning and modulation capabilities, phase continuous frequency transitions, high dynamic range, and good phase noise and stability. This set of characteristics is unique, and not easily realized with any other synthesizer architecture. In terms of price and power dissipation, components that accommodate real world applications have been available for only a few years. The availability of complete digital solutions and higher performance DACs, optimized for price and performance, is sure to drive this frequency synthesis option into many applications.

There are, of course, tradeoffs and limitations that designers must make to take advantage of this versatile synthesizer architecture. Applications requiring a high SFDR will find their options limited to the higher resolution DACs with lower clock rates and analog bandwidths. Applications that can take advantage of the higher speed DDS solutions may find the cost and power prohibitive. Designers must evaluate their system's architecture to determine if DDS offers an advantage over traditional frequency synthesis techniques, or if a modified architecture could take advantage of DDS and improve system performance. Applications which can usually take advantage of DDS include (but are not limited to) advanced military radar, high performance instrumentation, digital communications links, and commercial cellular communications.

CHOOSING DACS FOR DDS

Since DDS provides a frequency reference, it makes sense that its key specifications are in the frequency domain. While the phase noise and stability requirements may be directly related to specifications of the frequency reference, f_c, it is difficult to relate frequency range and SFDR specifications directly to high speed DACs that are suitable for DDS applications. High speed DACs are traditionally specified in the time domain. This is reasonable, since the traditional applications (video, analog-to-digital converter building blocks, fast tuning voltage references) are concerned with the time domain characteristics of the DAC. These specifications are constant without regard to the end application.

Earlier sections of this paper described the high performance characteristics of the DDS architecture and how its use is likely to grow in communication, instrumentation, and military applications. It would seem natural, with such an opportunity, that DAC manufacturers would design and specify devices especially for DDS applications and solve the problem of choosing DACs for DDS. To some degree, this is already happening.

Highly integrated DACs (including on-board registers, reference circuits, and reference amplifiers) optimized for waveform synthesis are available today, and manufacturers are providing applications support for their use in direct digital synthesizers. The problem of specifying the DAC is not solved quite so easily.

There are two reasons this is unlikely to change. The first has been explained, and is simply because various system applications of DDS have a wide range of dynamic requirements. The second reason involves the inconsistency of DAC performance over a wide frequency range. It is the combination of these two characteristics that makes it so difficult to specify the devices. If the applications requirements were narrowly focused, DAC manufacturers would simply specify the devices accordingly, and ignore the inconsistency in performance.

Alternatively, if the DAC's performance were to degrade in some predictable fashion (as an amplifier might) the manufacturer could supply data or a set of performance curves that would allow designers to predetermine system performance. As it stands, the manufacturer does well to indicate simply that the part is intended for use in DDS applications, and to supply some sample data to back up this claim.

Since high speed DAC manufacturers are often unable to specify the devices adequately, the task of deciding whether a part is suitable for a particular DDS system application is left to the system designer. It is the goal of the remainder of this paper to explain the relevant dc and time domain specifications of the DAC, and how they can be related to frequency domain performance. It will become clear that these specifications will not provide all the necessary information to predict the frequency domain performance; consequently, some recommendations on how to characterize the DACs for DDS applications are also presented.

DAC DC SPECIFICATIONS

Figure 7 illustrates both an ideal and real world transfer function for a 3-bit DAC. Manufacturers typically specify offset, gain error (full scale in relation to reference),

differential nonlinearity (DNL), and integral nonlinearity (INL) as an approximation of this transfer function. Output offset is usually defined as a constant dc offset in the transfer curve, and therefore has no effect on the frequency domain characteristics of the output. Gain defines the full-scale output of the converter in relation to its reference circuit. The gain error is usually specified tightly enough so there is little concern about its effect on the frequency domain performance.
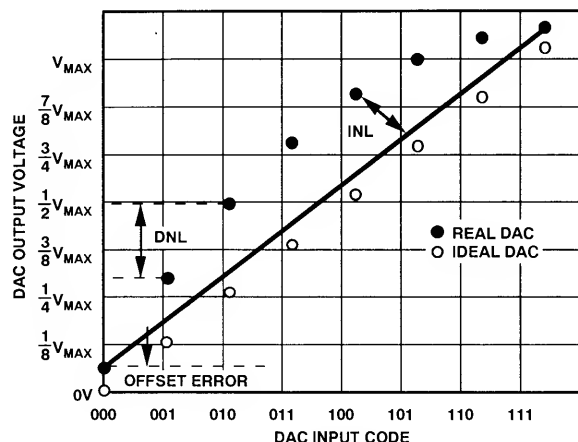

Figure 7. DAC Transfer Function

The linearity specifications, differential nonlinearity (DNL) and integral nonlinearity (INL), provide useful information about frequency domain performance. DNL is typically measured in LSBs as the worst case deviation from an ideal LSB step between adjacent code transitions, and can be a negative (less than 1 LSB step) or positive (greater than 1 LSB step) error. DACs with a DNL specification of <−1 LSB are not guaranteed to be monotonic. Figure 7 illustrates both positive and negative DNL errors, although the part is monotonic.

Armed with the DNL error of each code transition, a diligent individual could predict the frequency domain performance of an individual DAC's transfer function. This process would consist of modeling the digital part of a DDS; using the modeled DDS data as input to the DAC transfer function; collecting samples of the simulated DAC output levels, and performing a Fourier transform to observe frequency domain performance. This process would have to be repeated for every ratio of clock to analog output frequency. While this might be an interesting exercise, it is not a realistic way of selecting a DAC.

Manufacturers usually measure DNL for each code transition, but specify only a worst case error. For some applications (analog to digital converter applications, for example) this is sufficient to predict the errors contributed by the DAC's transfer function. For frequency synthesis, however, it raises some interesting questions. For example, all other things equal, will a DAC with a 1/2 LSB DNL error outperform a DAC with a 1 LSB error in the frequency domain?

The answer is not obvious. Theory predicts that a perfect DAC generating a full-scale sine wave will have an rms signal to rms noise ratio of 6.02N − 1.76 dB[1] over the full Nyquist bandwidth [ignoring (sine x)/x rolloff[2] where N is the resolution of DAC]. A perfect 12-bit DAC could therefore provide 70.48 dB of SNR. A measurement of the DNL error would predict a maximum reduction in this performance; for example a 12-bit DAC with a 1 LSB DNL error (every other quantization level missing) would have a minimum SNR of 64.46 [6.02 × (11) − 1.76] dB (if DNL were its only nonlinearity).

This reduction in SNR, however, does not predict how the distortion is spread over the Nyquist bandwidth. If the DNL errors caused the additional noise to be spread evenly over the Nyquist band, their effect would be negligible. However, if the DNL errors concentrate portions of the noise into a single frequency (usually a harmonic) this could limit the SFDR.

Getting back to the example, if the 1/2 LSB device has its maximum DNL error on a high percentage of its code transitions, and the 1 LSB device were perfect except for a single code transition, then the 1 LSB part would have better frequency domain characteristics. It is obvious from this example that DNL specifications can be misleading if not interpreted properly.

INL is measured as the worst deviation from a straight line approximation to the DAC's transfer function. The straight line approximation eliminates the dc errors (gain and offset) which have already been discounted from any frequency domain effect. In Figure 7, this straight line is drawn between the two end points. The INL characteristics of some devices are measured against a "best fit" straight line through the converter's code transitions. Like the DNL specification, the INL measurement is a worst case deviation. It does not indicate how many DAC codes reach this deviation, nor which direction away from the best straight line the deviation occurred.

Figure 8 illustrates how this specification might be misinterpreted. Each of the illustrated curves represents a transfer function that would have the same INL measurement, but three distinct effects in the frequency domain. For example, the transfer function of the "bow" INL curve will introduce a prominent 2nd harmonic distortion, while the symmetrical "S-curve" will tend to introduce 3rd harmonic distortion.
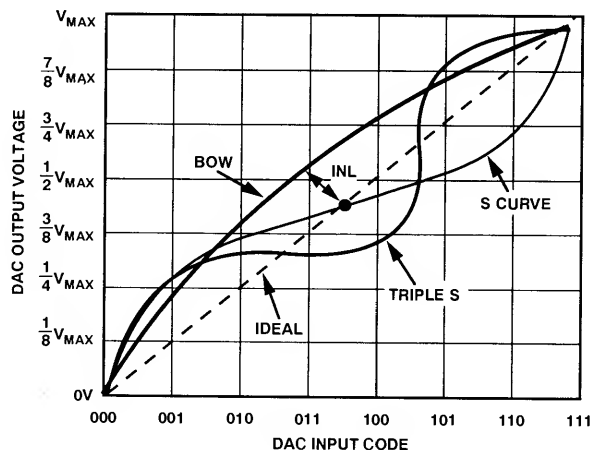

Figure 8. INL Curves—All 1 LSB

Most IC DAC designs recommended for use in frequency synthesis exhibit some predictable linearity pattern, with DNL and INL errors occurring due to architectural trade-offs and matching of process parameters. This pattern determines how the DNL and INL will contribute to the overall frequency content of the signal. Because the linearity patterns of any two specific DAC designs will be different, it is impossible to compare their frequency domain performance by comparing their linearity specifications.

Since deviations in DNL and INL will probably not affect the ac characteristics (glitch, slew rate, etc.) of the DAC, it is possible to characterize the effects of linearity on frequency domain performance for a specific DAC. This characterization would consist of selecting units that exhibit good DNL and INL characteristics, and comparing their frequency domain performance with parts that exhibit poor linearity. This allows designers to determine if tighter linearity grades of a specific DAC will improve a synthesizer's performance.

AC TIME DOMAIN DAC SPECIFICATIONS

The ac time domain specifications of a DAC relate to its code-to-code transitions as illustrated in Figure 9. It is difficult to predict the frequency domain effects of any of these specifications, and practically impossible to predict the frequency domain effects of their combinations.

It is, however, possible to consider what contributes to each one of these nonlinearities, and how its effect can be minimized in a final application.
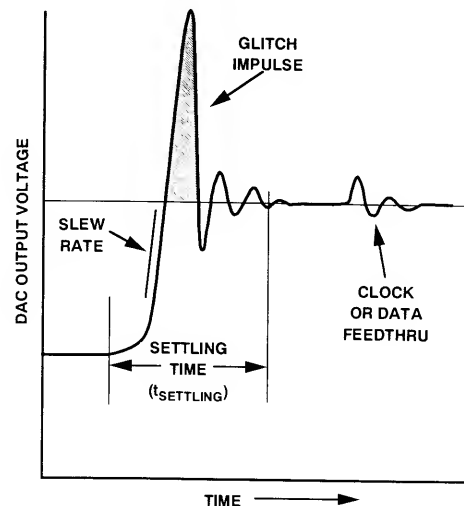

Figure 9. DAC Transition

The effect of DAC output slew rate on its frequency domain performance is not as straightforward as it is in linear applications due to the presence of glitch impulse and settling effects (discussed below). In general, a DAC with a high slew rate will produce an output transition that is closer to ideal than one which is slower. Designers should be careful to consider also the difference in rising and falling slew rate, as this will tend to concentrate energy in the harmonics of the fundamental output frequency.

Glitch impulse, often considered a key figure of merit in DDS applications, is simply a measure of the initial transient response (overshoot) of the DAC between two output levels. It is usually measured as the area of the transient, and ranges from 15 to 100 pV-s for devices commonly used in DDS applications. This transient is commonly associated with the time skew between the data bit transitions or by unequal propagation delays through the internal logic. In either case, the time skew would cause the DAC's output to approach an intermediate state, and probably add unwanted energy to the output frequency spectrum.

DAC architecture has a strong effect on the magnitude of the glitch impulse. Architectures employing an on-board register for data deskew and propagation delay matching, along with a segmentation of the major bits, have the lowest glitch impulse. The glitch illustrated in Figure 9 is a peak glitch interpretation. Designers are warned that not all manufacturers define glitch impulse in the same fashion, and comparison of data specifications is often misleading.

Before the advent of monolithic devices which minimize glitch impulse and other DAC nonlinearities, it was common practice to "deglitch" the outputs of hybrid converters using a track-and-hold amplifier (T/H). The idea is to place the T/H in "hold" mode during the DAC's transition to a new value. After the converter's output is settled to its final value, the T/H is allowed to return to "track" mode; and the nonlinearities of the T/H's transition to the newly acquired DAC output become the limit to ac performance.

T/H technology has been predominantly hybrid technology, making any possible improvement in performance a costly venture. Although the transition characteristics of the T/H may imply that deglitching could improve the DAC, overall performance may actually degrade because of limits in other performance characteristics such as slew rate and absolute accuracy. More recent monolithic T/H product offerings such as the AD9101 125 MSPS sampling amplifier may offer marginal improvement in overall ac performance at a reasonable cost.

After the initial transient, the DAC's output will settle to its final value. The settling time is usually measured as the time from the digital inputs' transitions to the time the DAC's output settles to within a certain error band (usually 1/2 or 1 LSB) of the final output value. Many manufacturers understandably argue that the digital propagation of the converter should not be included in this specification. Settling time should instead be measured as the interval from the time when the DAC's output leaves the error band around its initial value, and when it settles within the error band around its final value. This definition more accurately describes the non-ideal characteristics of the code transition.

Slew rate, glitch impulse, and settling time characteristics of a DAC's output are all dependent on the output

loading circuit. Stray capacitive loading generally increases all of them. Most high speed DACs are current output devices and require an external load resistor to produce an output voltage reference. Design of the final DDS circuit should pay close attention to the DAC load, with signal traces kept short as possible, and/or with impedance matching techniques being used.

Feedthrough of the clock or data transitions to the DAC's output also adds to the frequency content of the DDS's output spectrum. These effects are often related to the test circuit layout and can be minimized with good circuit layout techniques. Many DAC manufacturers suggest the use of a series resistor in the input data connections to minimize data feedthrough. This series resistance works with the input capacitance of the DAC to form a low pass filter, and may alter the setup and hold characteristics of the device. This technique is not recommended on the clock connection as it will tend to add jitter (phase noise).

Theoretically relating the ac time domain characteristics of a DAC to its frequency domain performance is not feasible, nor recommended. One simple observation about the ac nonlinearities is possible: they become a more significant portion of each clock cycle as the clock or analog output frequency of a DDS system increases. Consequently, one can expect the spurious frequency content of the DAC to degrade at higher frequencies. Many DDS designs will operate well below the maximum clock frequency of the DAC to get optimum performance, and translate the output to a higher frequency using analog techniques (mixers).
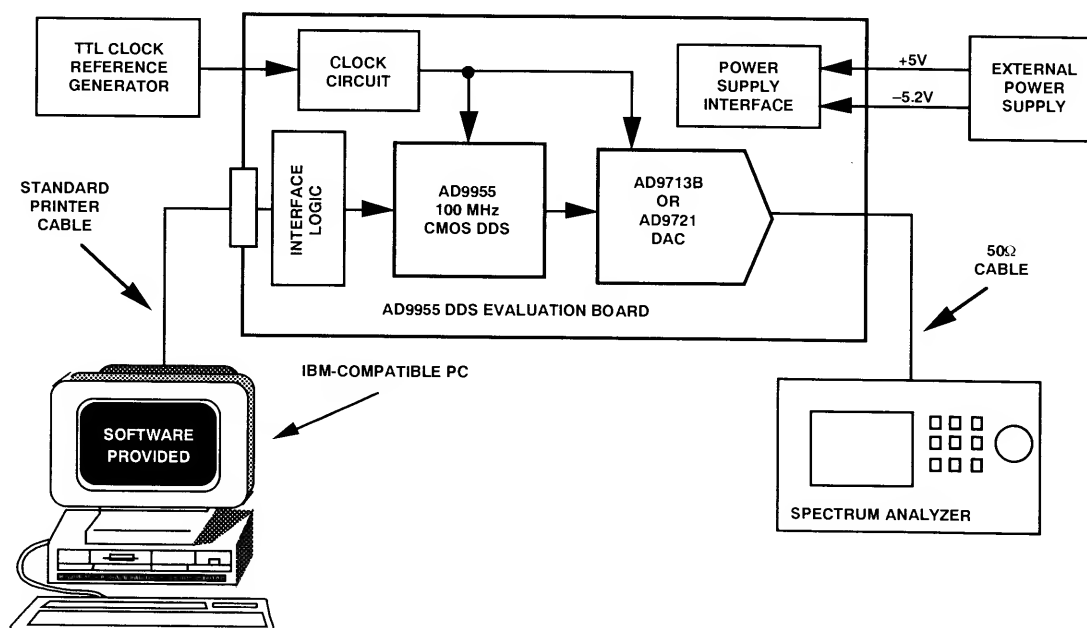

Figure 10. AD9955/PCB DDS Evaluation Board Setup

This degradation at higher clock rates, however, may not be the dominant effect in output spectral purity. Instead, for any given bandwidth of interest, the frequency domain performance may vary as the ratio of the clock to analog frequency changes. Spurious distortion will tend to appear in the output spectrum at alias frequencies, $Af_a \pm Bf_c$, where A and B are integers. The effects of these aliased components are concentrated near the fundamental frequency when the fundamental output frequency is nearly an integer fraction of the clock frequency. Designs that cover wide analog output bandwidths may find it hard to avoid this effect even at low clock rates, while narrow bandwidth applications may be able to find a particularly "clean" segment of output spectrum that allows taking advantage of the higher clock rates. This would require careful characterization of the DAC for the bandwidth of interest.

While the ac specifications of the DAC will not allow accurate predictions of the output spectral purity, these specifications are recommended as guidelines for determining if the DAC is worthy of being characterized for use in a DDS application. Devices that exhibit high slew rate, low glitch and feedthrough, and fast settling times should be considered first. Once the decision has been made to characterize, these time domain specifications are also useful as a guideline to determine if the DAC has been optimized in the test fixture. A block diagram of how this test is performed is illustrated in Figure 10.

The power supplies should be the linear (nonswitching) type if possible, and appropriate filtering needs to be included in the circuit to reduce noise to a minimum. For the purposes of evaluation, separate digital and analog supplies should be used as recommended by the DAC manufacturer. Once the characterization is complete, the designer can explore how combining digital and analog supplies would affect the spectral purity of the DDS.

The external controlling circuits should allow complete access to all digital circuitry. This may not seem appropriate if all of the features are not required in the final application, but may simplify the characterization process and allow the test set to be used in future design phases. The clock distribution is usually generated by an instrument during characterization, allowing a wide range of reference frequencies. The clock distribution circuit serves to buffer the incoming reference and sets up the appropriate timing between the digital part of the DDS and the DAC. The important specifications here are the DAC's input data setup and hold times. Even a slight violation of these specifications can cause an increase in glitch impulse if the latch is transparent, or capture of erroneous data if the latch is edge triggered.

Once the test set is constructed, functionality must be verified. This is most easily accomplished with a relatively slow clock rate. Checking functionality includes investigating the stability and level of the power supplies (at many points in the circuit), operation of the digital section of the DDS, stability of the DAC reference circuit, and the timing relationships at the DAC interface. The output of the DAC should be analyzed in the time domain to verify that the DAC is meeting the ac characteristics specified in the data sheet. If they are significantly different, the differences must be resolved.

After the circuit is functioning properly, the DAC can be characterized for its frequency domain performance. There are two approaches that can be taken: to characterize the part in general, or to characterize for a particular application.

To characterize the part for general use, a matrix of conditions is suggested. A maximum usable clock rate can be identified from the data sheet. The converter should be tested at several clock frequencies over the entire range. For example, a DAC billed as a 50 MSPS converter might be tested for clock rates of 2, 5, 10, 20, 35, and 50 MSPS. At each of these clock rates, the SFDR of the output should be characterized using a spectrum analyzer for several clock to analog output frequency ratios.

Determination of these ratios is a subject of debate, but the following selections would give a decent understanding of performance:

1. A low ratio of $<1/10\ f_c$, to determine the performance for the given clock frequency, and

2. Ratios of $1/3\ f_c \pm \Delta f$ and $1/4\ f_c \pm \Delta f$ to analyze the level of spurious frequency energy at alias frequencies.

For each of the three analog selections, the characterization of the spectrum will consist of two parts. The first considers performance across the entire Nyquist band, or a reasonable section of it. The second considers a narrow band around the fundamental to look for unfilterable spurious distortion. This process can become quite time consuming because of the time required for broadband, high dynamic range sweeps of the analyzer. A sample of data characterizing a prototype of the AD9721, a 10-bit, 100 MSPS TTL converter is shown in Figure 11.
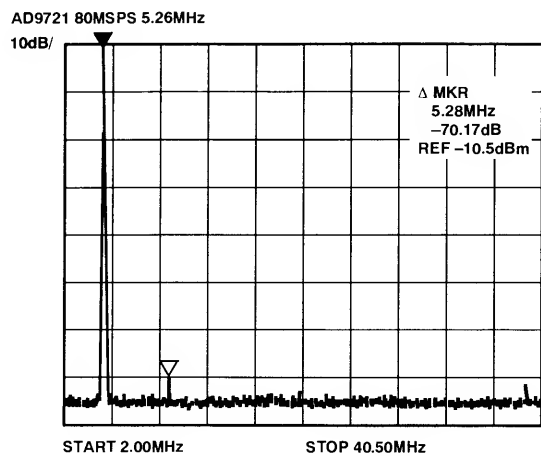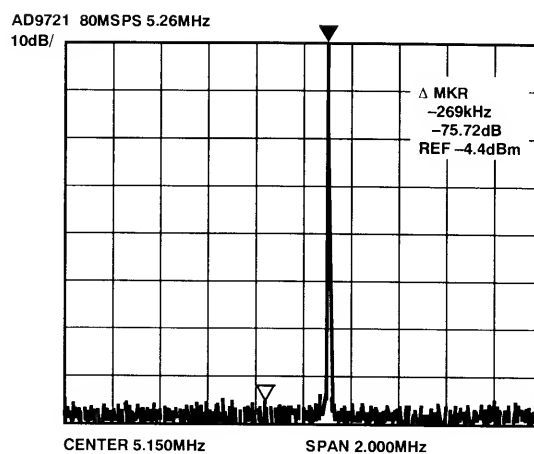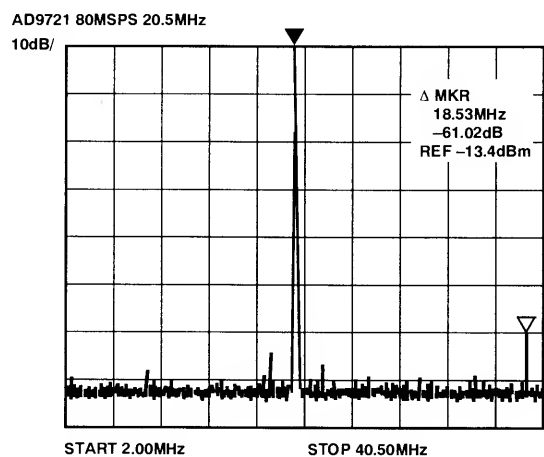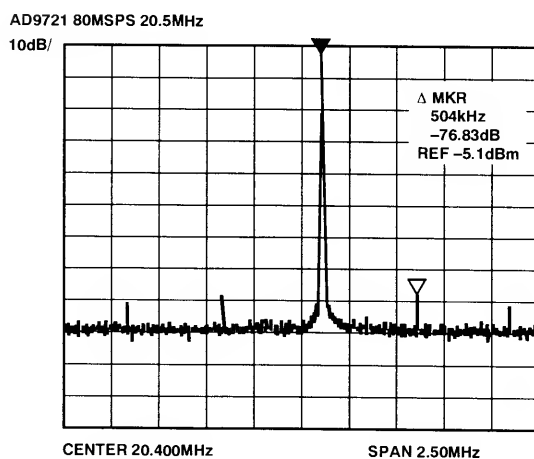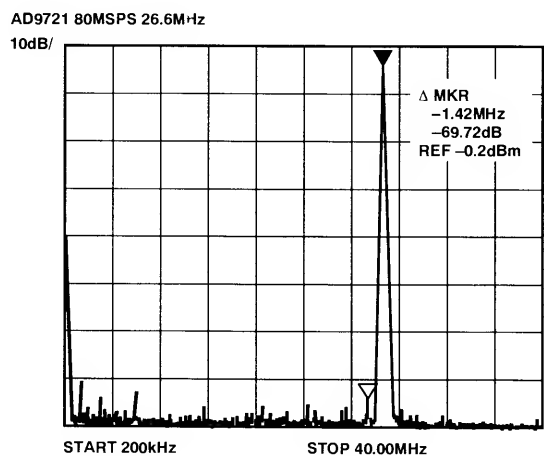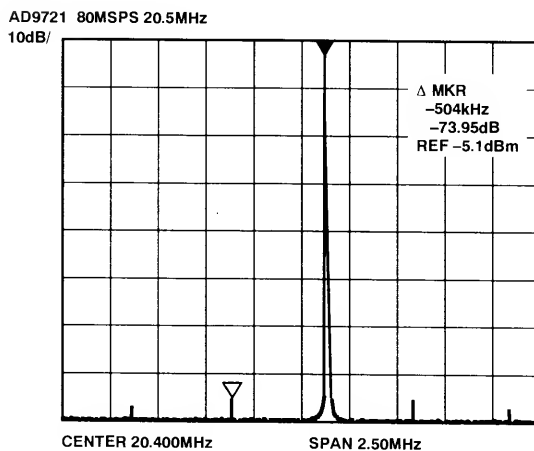
Figure 11. DAC Characterization Plots

Once this general characterization is complete, some limited conclusions about the DAC can be made. For instance, in the worst case, the part may exhibit poor performance across the test matrix; and the device could be disqualified from any further consideration. Encouraging results might indicate that the device provides enough SFDR for the application of interest.

If so, the application specific characterization process begins. This would consist of developing a frequency plan to use the DAC to generate an IF band to be used in the system. Characterization of the DAC would focus on determining the clock rate which could give the best performance in this band.

CONCLUSION

This application note has explored the architecture and advantages of DDS, and pointed out that the performance-limiting block is the DAC. It should be clear from the discussion that frequency domain performance cannot be accurately predicted from device specifications such as linearity, glitch impulse, slew rate, and settling time. Instead, testing is required to characterize the spurious free dynamic range of a converter for a range of clock and analog frequencies. A suggested method of characterizing the DAC was also discussed.

REFERENCES

1. Seminar Notes: "Fundamentals of Frequency Synthesizer Design," Instructor: Fred Studenberg. Nov. 6–8, 1990, University of Maryland Center for Professional Development.

2. W. R. Bennett, "Spectra of Quantized Signals." BSTJ 27, pp 446–472, July, 1948.